

Implementation of a Deep Space Receiver on 350K gate GaAs gate arrays.

Gary R. Burke, Terrence W. Chow, J. Scott Graham,
James E. Kowalski, William D. Whitaker, Robert A. Johnson

Jet Propulsion Laboratory,
M/S 238-737, 4800 oak Grove Drive
Pasadena CA91 109

Abstract

A set of 3 GaAs ASICs have been designed which together form part of the BLOCK V Digital Receiver. Each ASIC contains approximately 150K-1 70K used gates. This paper describes the design methodology for the 3 GaAs ASICs, which ensures successful timing, testability, and functionality.

Introduction

The recent availability of high density, high speed GaAs gate arrays, is making a major impact on the design of digital signal processing systems. Traditional analog functions can be put into digital logic, with considerable improvement in stability and accuracy. In this paper we discuss the design of a set of ASICs for a deep space digital receiver (BLOCK V). The ASICs use a purely digital approach. Such functions as oscillators, mixers and phase detectors are implemented by a combination of multipliers, adders and registers [2]. The large number of gates required (about 160K per ASIC) and the high speed requirement (160 MHz), would have been difficult to implement in an alternative technology. CMOS would have been too slow, and ECL did not supply the gate density required, and consumed higher power. A custom design approach was ruled out, since the schedule requirements were too stringent.

The design is partitioned into five ASICs (Fig. 1), with three different ASIC types. The first ASIC (Jetcar) performs the carrier demodulation. It contains the carrier demodulator and the carrier oscillator, as well as other functional blocks[3]. The sampled IF signal is multiplied by 16-bit sine and cosine signals at 160 MHz, and the products are lowpass filtered and decimated to 80 MHz samples. The Jetsub ASIC performs the




Figure 1 ASIC chip set

subcarrier demodulation. Subcarrier demodulation mixes both inphase and quadrature signals against sine and square wave subcarrier references. Subcarrier loop accumulators form summations for subcarrier phase tracking, and lock indication. The Jetsym chip processes the 'symbols', which are the extracted data bits. Symbol loop accumulators form summations for symbol phase tracking, lock indication, symbol signal to noise ratio, and symbol count over the accumulation interval.

The ASICs each have a VME bus interface. These allow data to be read and written to registers on the ASICs, from a micro-controller (on another board). The interface levels for this connection are TTL, but the high speed interconnections are all ECL. In this design we mixed both TTL and ECL 1/0s.

Technology

All three ASICs were designed using the Vitesse FX 350KB gate array [1]. These are channelless gate arrays, with 809 rows and 216 columns. There are 600 pads on the die, of which 108 are input pads, 218 are input/output pads and the remainder are VTT power (-2 volts), VTTL power (+3.3 volts) and Ground. The chip is packaged in a 557 pin cavity-down PGA, with a heat spreader. Both ECL and TTL level I/Os are available.

The Vitesse gate arrays are implemented with H-GaAs III, which is a 0.6 μ m process with four levels of metal interconnect. Three of these levels can be used for signal interconnect (routing): Metal 1, Metal 2 and Metal 3. The personalization of the Gate Arrays consists of patterning of seven mask layers: Via 1, Metal 1, via 2, Metal 2, Via 3, Metal 3 and Via 4. Metal 1 is used for routing and for interconnect inside the core and 10 macros. Metal 2 is used exclusively for routing and Metal 3 is used for both power and routing. Metal 4, a non-programmable layer, is used for power distribution. Figure 2 shows the Jetsym chip.

A cell occupies 1 row and 0.5 columns. All logic gates occupy multiple cells. A simple unbuffered inverter takes 2 cells, a 2-input buffered NOR gate occupies 4 cells. Although various floor plans are available, we used a 67% floor plan, which means every 3rd column is used for routing only. Additionally, we did not exceed a utilization of 75% of the remaining cells. We did not experience any routability problems with this plan.

The basic gate structure is a DCFL NOR gate, with a depletion load. The DCFL logic switches around the enhancement mode threshold, which is about 250 mV above the Source Voltage, VTT. A typical gate delay is 115 ps at 0.2 mW for an unbuffered 2 input NOR with a fanout load of 1 and a wire length of 0.21 mm. A buffered 2 input NOR with a fanout of 3 and 0.63 mm wire has a gate delay of 130 ps at 0.7 mW.

The Vitesse gates are available in different drive strengths; high power drivers are used to drive longer lines, while low power gates are used to save chip area for local interconnects.

Timing

Using a gate array for high performance design presents some problems, since the design is very much influenced by the gate speed. The position and number of pipeline stages is dependent




Figure 2 Jetsym chip

on the eventual gate delay, and the amount of buffering is dependent on the eventual wire lengths. The delay through a DCFL GaAs gate is very dependent on the capacitance on the output of the gate, largely resulting from interconnect wiring. The wire length is of course very dependent on placement. Unlike a custom chip, there is little control over the final placement. Although pre-placement of the array can be used, this is impractical for large designs. In consequence, the pre-layout design depends on knowing the interconnect wiring lengths, which are not known until after auto-layout. To avoid iterating through design and layout, we used a new approach which we call IFA (intelligent front annotation). This generates a worst case estimate of the wiring lengths, and also generates timing constraints for place and route. Using this approach, no iterations of place and route were necessary to correct timing problems.

The chips are designed hierarchically. At the top level is the I/O ring and the core logic. Within the core logic are functional blocks (e.g. Carrier Demodulator). These functional blocks are designated regions on the chip in floor-planning. I/O locations are defined, and regions assigned to assist the data flow through the chip. Large blocks within these regions (e.g. a multiplier) are assigned to soft-groups. The soft-groups are formless, and can be placed anywhere within a region. They are used to direct the place and route software as to which macros should be in near proximity. Each net is assigned a maximum length (MAXLENGTH). The place and route iterates to meet these constraints.

The MAXLENGTH value may be assigned by the designer or automatically by IFA.

Traditionally, Front Annotation (FA) has been used in ASIC designs, to give an estimate of wire delay before placement. The FA algorithm uses only the fanout on the net to estimate wire length. This is a very poor estimate, since a high fanout net could be constrained in one area of the chip, while a low fanout path may stretch over a long distance. IFA uses the hierarchy of the design, and the known region information to provide a realistic estimate of wire length. To estimate the distance between any two nodes, the common hierarchical block is found. The estimated size and aspect ratio of this block is used to estimate the wire length. Other taps on the same net and within the same block are given a smaller estimate, and the wire sizes become smaller as the fanout is increased,

Factor	Jetcar	Jetsub	Jetsym
O- 1 X	14356	13506	15748
1-2 X	178	89	176
2 - 3 X	6	25	60
3 - 4 x	3	2	26
4-5x	0	0	1
5-6X	2	0	0
6-7X	0	0	0

Figure 3 Histogram of Wire lengths

The wire estimates from IFA are used for static timing analysis. This ensures that there are no path delays which exceed the specified limit. The wire estimates are also used to constrain the place and route algorithm. The final Back Annotated wire lengths are then compared to the estimated values to see whether the constraints have been met. In all cases the constraints were met in the first 'real' place and route run. Figure 3 shows the results after placement for all 3 chips, showing nearly 99% of the nets were under 1 X the constraint value. Back-annotated timing exceeded the specification. There was no need to adjust the placement to improve the timing.

Testing

A set of functional vectors was generated. This was fault graded. Undetected faults were input to an Automatic Test Pattern Generator (ATPG) tool. This increased fault coverage to acceptable levels (>97%). To ensure testability, individual blocks were run through ATPG during the design

process and testability problems were corrected. To assist in generating high-coverage test vectors, internal registers were connected to a scan path. The number of scannable flip-flops was very large (2000-2400 bits). Rather than create one huge scan path, it was divided into multiple scan paths of up to 128 bits each. This does not affect fault coverage, but does considerably reduce the number of test vectors required,

During wafer probe, the impedance of the power connections is not as low as in the packaged part, and the noise due to simultaneous switching outputs (SSO) can affect the testing. We took two approaches to avoid this problem. Firstly, a parametric tree was created, whereby every input can be tested individually, for VI]] and VIL without any SSO problems. Secondly, the outputs were arranged to be switched off in groups, using a dedicated set of inputs. For example there are 9 enable pins in Jetsym, so that only 1/9 of the outputs need be enabled at any one time. By running the same vector set 9 times, with a different output set enabled each time, the SSO problems are completely avoided.

The tester will not run as fast as the ASICs, so that the parts are not tested at speed. To ensure that the parts are not slow, 3 delay paths are used in each ASIC. The delay through these combinatorial paths is compared with the estimated delay. Because of the variation in propagation delay on the chip (25%), the path delay may not represent the actual device speed. To reduce this possibility, very large paths were taken, so that the delay paths would cover most of the chip, and experience all of the possible variations. One path was the parametric tree. Another path was a set of invertors, which was put into every region to form a long chain. A Built-in-Self-Test circuit is used to test the parts at speed in the system.

Design issues

For fast two's complement multiplication, the Baugh-Wooley algorithm was implemented using a Wallace tree of carry save adders for column compression. The partial products are formed in subblocks of 16 NOR gates, grouped to form 4x4 arrays of common input loads. The partial products are fed to column compression subblocks consisting of adder trees for 4-bit groups. IFA generated constraints from this partitioning were sufficient for intrablock connections. Maximum wire length constraints were forced to 2000µm for the least significant inter-column compression group connections. Although unbuffered gates are available in Vitesse's 350K library, only 3 were used in this

multiplier. Capacitance loading of the unbuffered adders would have made the adder too slow. All 101 adders were the 1X drive cells; no 2X drive adders were needed. Eleven half adders were also used. Data arriving at the unit is distributed to the partial product NOR gates, and then propagates through three carry save adders (CSA) before the pipeline register. Next the data travels through one more CSA and then the carry lookahead adder (CLA) to the product register. Placing the pipeline stage within the Wallace tree used 57 flipflops rather than 45 which would be needed at the CLA inputs. The 4-bit adder macrocell was used for most of the final adder, for efficient use of chip area.

Because of the large die size (15mm~ side), a large clock skew was expected between different parts of the chip. The high speed of the GaAs flip-flops would make hold time problematical. To avoid this, we used two techniques. The majority of the flip-flops in the core of the ASIC were connected to a fixed clock tree. The clock tree is predefined for the masterslice, and has guaranteed low skew. It uses clock drivers assigned to the first 5 and the last 5 rows, as well as a 3X2 area in the middle of the chip. For the boundary scan logic, which operates at a different and low speed clock, we connected a series of clock buffers in opposite order to the boundary scan data flow. Groups of 8 boundary scan registers share a clock buffer. The clock buffers are preplaced, so there is no chance of a hold time problem.

Chip-to-chip timing is a problem due to the variation in speed between devices. Large amounts of data transferred between Jetsub and Jetsym, are generated and sampled by the same clock edge. Since the chips are physically in close proximity, transmission delays are small. A hold time problem would occur when the Jetsub part is fast enough to change the data, as it is being clocked into Jetsym. To avoid this, additional delay was inserted into each of the critical Jetsym inputs. To ensure consistency in delay values, without preplacement, a circuit was used which achieved the majority of its delay (1 ns) by gate capacitance rather than wire capacitance. The circuit was also designed to provide a symmetrical delay time.

We used high power driver macros to drive long nets between regions, because of the large capacitive load. This can cause an overdrive problem. The high power driver will put too much current into the receiving enhancement mode FET Gate-Source diode, which results in a poor VOL level for that transistor. To avoid this problem we used a 'LOAD' gate which puts a dummy load on the net. However, although it is convenient to put the LOAD at the source, this causes another

problem: Long nets may be subject to 'current hogging' due to variations in VTT across the die. The current will tend to flow in the receiving FET with the lowest VTT, causing a possible overdrive in that device, and causing an under-drive problem in the other load(s). This will be detected by an ERC program, but not until after place and route. To avoid this problem, the LOAD cells were placed at the destination region, rather than the source region.

Conclusion

High speed, large scale digital GaAs ASICs were successfully designed. Care was taken that all paths would meet the speed requirement. Prototype versions of all 3 parts have been delivered. One part (Jetcar) has been extensively system tested, with excellent results. This part although designed for 160 MHz, performs correctly twice this clock speed. At the time of writing, testing is still proceeding on Jetsub and Jetsym; no significant errors have been found so far.

Acknowledgment

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The project would not have been possible without the assistance of many engineers at Vitesse Semiconductor, especially Tony Anderson, Patrick Van Eijk, and Robert Randall; the JPL BLOCK V ASIC group especially Jeff Buchanan, Barbara Lam, Fardad Siavoshi and Matt Thomson; the JPL BLOCK V Analysis Group including MiMi Aung, Jeff Berner, and James Layland; Ernest Stone for project management, and William Hurd for technical direction.

References

- [1] Vitesse Semiconductor, "FX/VIPER Gate Array Macrocell Library Release 4.2", March 1993
- [2] Gary R. Burke, "ASICs in Space" ASIC and EDA, pp. 20-29, March 1993.
- [3] S. Graham, M. Thompson, R. Johnson, "A GaAs ASIC Chip Set for NASA's Deep Space Network," 1992 IEEE ASIC Conference and Exhibit, pp. 57-63.

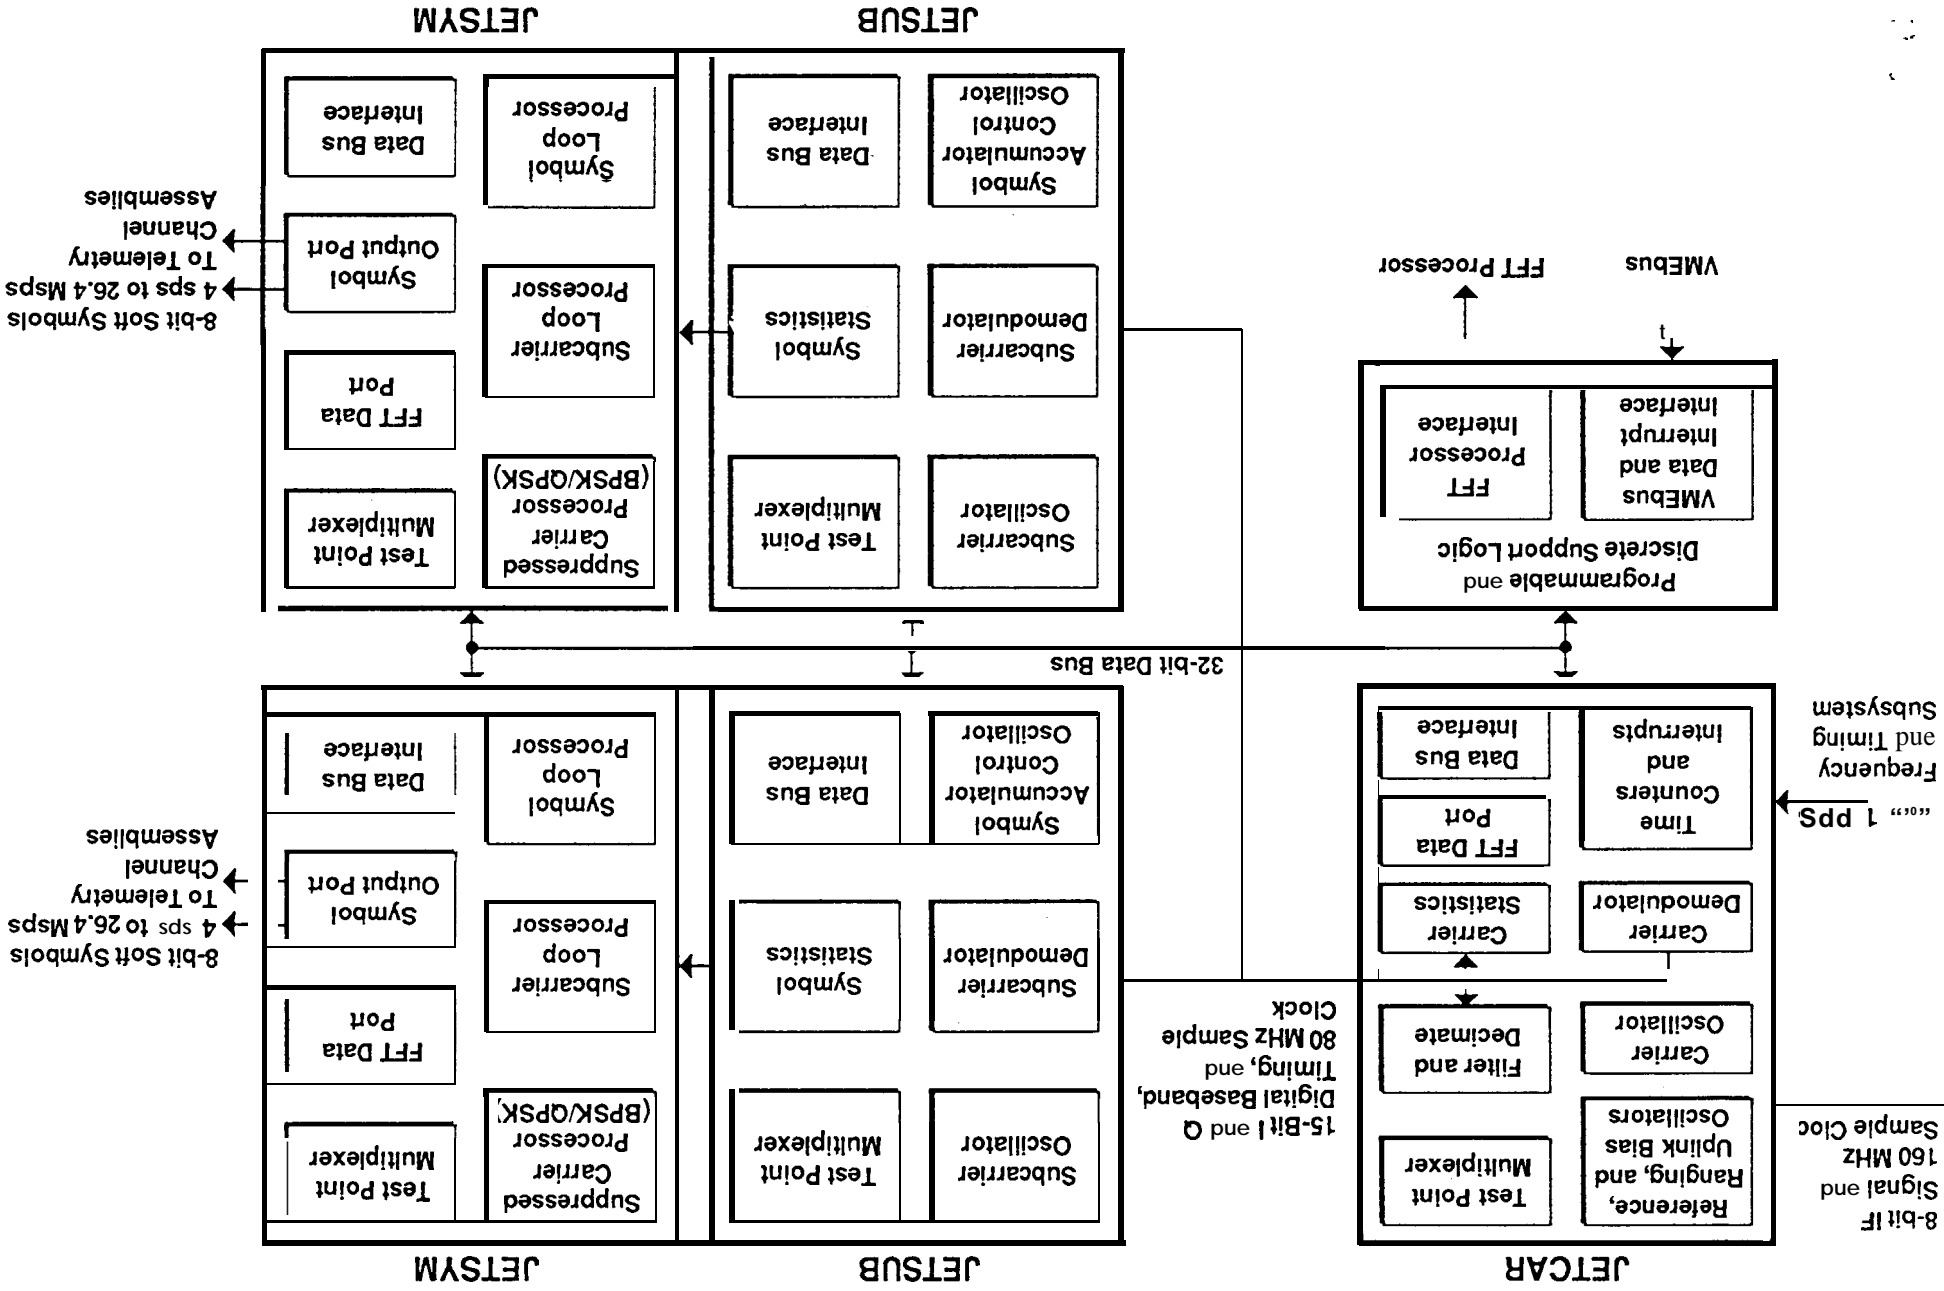
Spares

Testability was a concern since expected yields are low, requiring a high degree of fault coverage to avoid significant 'escapes'. We could not rely on easily detecting a fault. It may not show up easily in the system due to inherent fault tolerance. The receiver is designed to extract a signal buried in noise, a chip defect could in effect just add to this noise, in that case the receiver will still function, but not as well as it should,

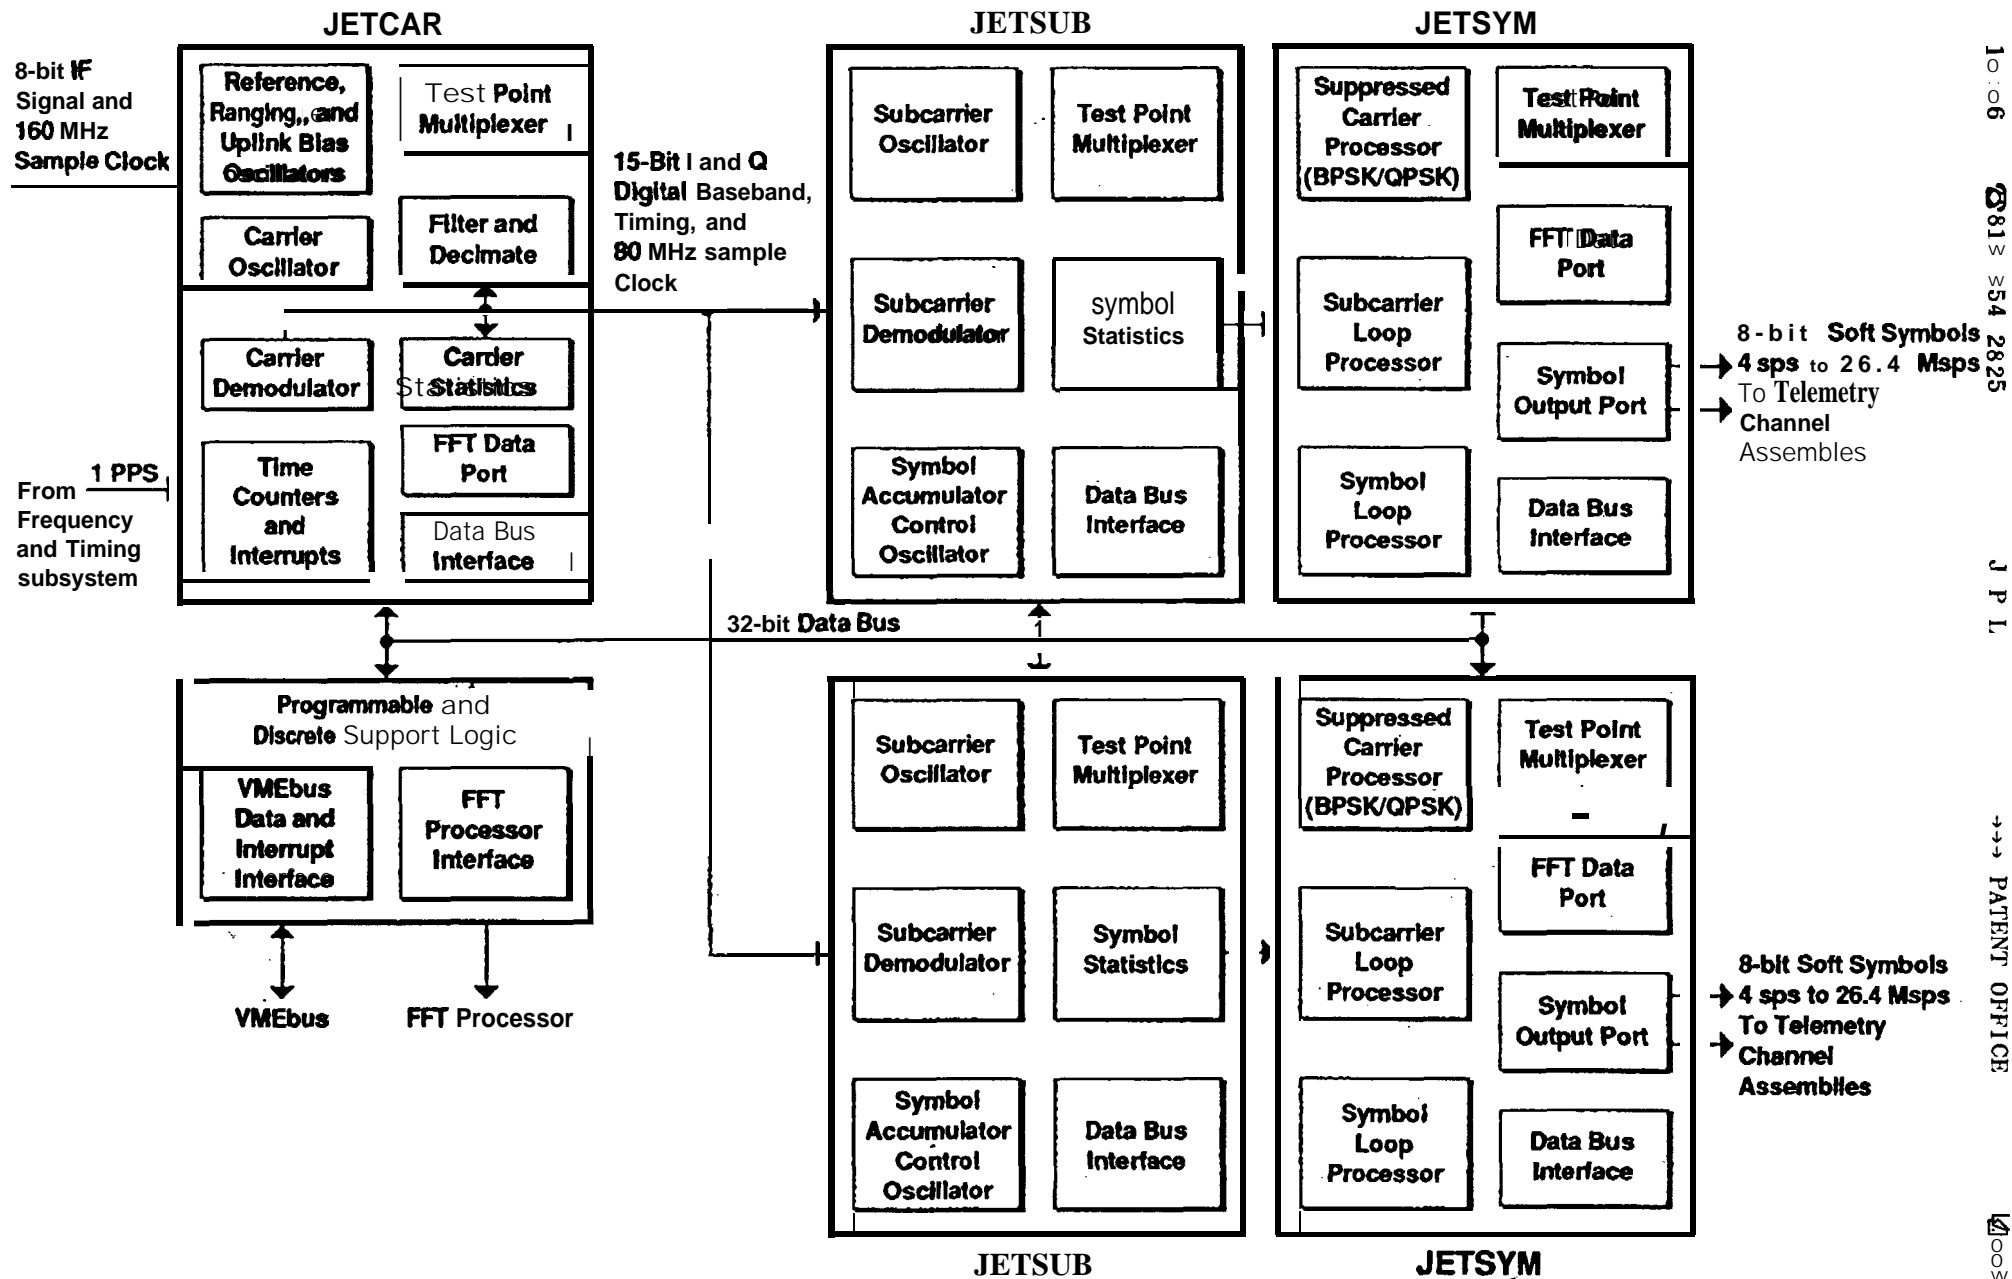
The carry lookahead macrocell was not used because no carry in was needed, allowing simpler implementation,

The fabrication process is completed by deposition of a protective passivation layer which leaves only the pad metallization exposed. The H-GaAs 111 process uses a total of 14 masks steps.

DIGITAL SIGNAL PROCESSOR ASIC CHIP SET



DIGITAL SIGNAL PROCESSOR ASIC CHIP SET






FIG 2